IN THE CLAIMS:

1. (Currently Amended) A method for uniformly distributing data transmitted by a server over a plurality of underlying links of an aggregate within a computer network, the method comprising the steps of:

defining a unit of data as a datagram having an Internet protocol (IP) identifier (ID);

apportioning ~~each~~ the datagram into at least one fragment at the server;

associating each fragment to an underlying link of the aggregate on the basis of ~~an Internet protocol (IP) identifier (ID) of each datagram~~ the IP ID and a number of active links of the aggregate; and

transmitting the fragment over its associated underlying link from the server to the computer network to transmit fragments of the datagram over the same associated underlying link.

2. (Original) The method of Claim 1 wherein the step of associating comprises the step of producing a result representing a remainder upon dividing the IP ID by the number of active links.

3. (Original) The method of Claim 2 wherein the step of associating further comprises the steps of:

calculating the IP ID of each datagram in a sequential manner; and

rotating the fragments of each datagram among all the underlying links to thereby ensure that all fragments having the same IP ID are provided to the same physical link of the aggregate.

4. (Original) The method of Claim 1 wherein the step of associating comprises the steps of:

logically combining the IP ID with a predetermined mask to produce a quantity;

right shifting the quantity a predetermined number of places; and

establishing a threshold at which a group of data is forwarded to each underlying link of the aggregate.

5. (Original) The method of Claim 4 wherein the step of associating further comprises the step of producing a result representing a remainder upon dividing the right shifted logically combined quantity IP ID and predetermined mask by the number of active links.

6. (Previously Presented) A method for uniformly distributing data transmitted by a server over a plurality of underlying links of an aggregate within a computer network, the method comprising the steps of:

defining a unit of data as a datagram;

apportioning each datagram into at least one fragment at the server;

associating each fragment to an underlying link of the aggregate on the basis of an Internet protocol (IP) identifier (ID) of each datagram and a number of active links of the aggregate, wherein the step of associating includes

logically combining the IP ID with a predetermined mask to produce a quantity,

right shifting the quantity a predetermined number of places,

establishing a threshold at which a group of data is forwarded to each underlying link of the aggregate, and

producing a result representing a remainder upon dividing the right shifted logically combined quantity IP ID and predetermined mask by the number of active links, wherein the IP ID is a 16-bit value, the predetermined mask is 0xFF80 and predetermined number of right shifted places is 7, and wherein the group of data comprises 128 IP IDs;

transmitting the fragment over its associated underlying link from the server to the computer network.

7. (Original) The method of Claim 6 wherein the group of data comprises one of 128 different transport control protocol (TCP) fragments and 128 different user datagram protocol (UDP) datagrams.

8. (Original) The method of Claim 7 wherein each UDP datagram comprises up to 23 fragments.

9. (Original) The method of Claim 1 further comprising the steps of:

loading at least one data buffer of the server with the at least one fragment;

fetching the fragment from the data buffer; and

loading at least one queue of the server with the fragment, the queue associated with the underlying link.

10. (Original) A system adapted to uniformly distributing data over a plurality of underlying links of an aggregate within a computer network, the system comprising:

a processor;

a memory coupled to the processor and having locations addressable by the processor;

an operating system resident in the memory locations and executed by the processor, the operating system configured to implement a modified load balancing technique that defines a unit of data as a datagram, the operating system comprising an Internet Protocol (IP) layer that apportions the datagram into at least one fragment, the operating system further comprising a virtual interface process that associates the fragment to an underlying link of the aggregate on the basis of an IP identifier (ID) of the datagram and a number of active links of the aggregate; and

at least one network adapter coupled to the memory and processor that cooperates with a network driver of the operating system to transmit the fragment over the associated underlying link to the computer network.

11. (Currently Amended) Apparatus for uniformly distributing data transmitted by a server over a plurality of underlying links of an aggregate within a computer network, the apparatus comprising:

means for defining a unit of data as a datagram, the datagram having an Internet protocol (IP) identifier (ID);

means for apportioning ~~each~~ the datagram into at least one fragment at the server;

means for associating each fragment to an underlying link of the aggregate on the basis of ~~an Internet protocol (IP) identifier (ID) of each datagram~~the IP ID and a number of active links of the aggregate; and

means for transmitting the fragment over its associated underlying link from the server to the computer network to transmit fragments of the datagram over the same associated underlying link.

12. (Original) The apparatus of Claim 11 wherein the means for associating comprises means for producing a result representing a remainder upon dividing the IP ID by the number of active links.

13. (Original) The apparatus of Claim 12 wherein the means for associating further comprises:

means for calculating the IP ID of each datagram in a sequential manner; and

means for rotating the fragments of each datagram among all the underlying links to thereby ensure that all fragments having the same IP ID are provided to the same physical link of the aggregate.

14. (Original) The apparatus of Claim 11 wherein the means for associating comprises:

means for logically combining the IP ID with a predetermined mask to produce a quantity;

means for right shifting the quantity a predetermined number of places; and

means for establishing a threshold at which a group of data is forwarded to each underlying link of the aggregate.

15. (Original) The apparatus of Claim 14 wherein the means for associating further comprises means for producing a result representing a remainder upon dividing the right shifted logically combined quantity IP ID and predetermined mask by the number of active links.

16. (Currently Amended) A computer readable medium containing executable program instructions for uniformly distributing data transmitted by a server over a plurality of underlying links of an aggregate within a computer network, the executable program instructions comprising program instructions for:

defining a unit of data as a datagram, the datagram having an Internet protocol (IP) identifier (ID);

apportioning ~~each~~ the datagram into at least one fragment at the server;

associating each fragment to an underlying link of the aggregate on the basis of ~~an Internet protocol (IP) identifier (ID) of each datagram~~the IP ID and a number of active links of the aggregate; and

transmitting the fragment over its associated underlying link from the server to the computer network.

17. (Original) The computer readable medium of Claim 16 wherein the program instruction for associating comprises a program instruction for producing a result representing a remainder upon dividing the IP ID by the number of active links.

18. (Original) The computer readable medium of Claim 17 wherein the program instruction for associating further comprises program instructions for:

calculating the IP ID of each datagram in a sequential manner; and

rotating the fragments of each datagram among all the underlying links to thereby ensure that all fragments having the same IP ID are provided to the same physical link of the aggregate.

19. (Original) The computer readable medium of Claim 16 wherein the program instruction for associating comprises program instructions for:

logically combining the IP ID with a predetermined mask to produce a quantity;

right shifting the quantity a predetermined number of places; and

establishing a threshold at which a group of data is forwarded to each underlying link of the aggregate.

20. (Original) The computer readable medium of Claim 19 wherein the program instruction for associating further comprises the program instruction for producing a result representing a remainder upon dividing the right shifted logically combined quantity IP ID and predetermined mask by the number of active links

21. (Currently Amended) A method for distributing data over a plurality of network links within a computer network, comprising the steps of:

defining a unit of data as a datagram, the datagram having an Internet protocol (IP) identifier (ID);

apportioning ~~each~~ the datagram into at least one fragment;

associating ~~each~~ the datagram to a network link of the plurality of network links according to a round robin policy based at least in part on ~~an Internet protocol (IP) identifier (ID) of each datagram~~ the IP ID;

transmitting the fragments of the datagram over the datagram's associated network link.

22. (Previously Presented) The method of claim 21 wherein the step of associating is further based, at least in part, on a number of network links in the plurality of links.

23. (Previously Presented) The method of claim 22 wherein the step of associating comprises the step of:

producing a result representing a remainder by dividing the IP ID by the number of network links.

24. (Previously Presented) The method of claim 23 wherein the step of associating further comprises the steps of:

calculating the IP ID of each datagram in a sequential manner; and

rotating the fragments of each datagram among all the network links of the plurality of network links to thereby ensure that all fragments having the same IP ID are provided to the same network link.

25. (Previously Presented) The method of claim 21 wherein the step of associating comprises the steps of:

logically combining the IP ID with a predetermined mask to produce a quantity;

right shifting the quantity by a predetermined number of places to create a Previously Presented quantity; and

establishing a threshold at which a group of data is forwarded to each network link of the plurality of network links.

26. (Previously Presented) The method of claim 25 wherein the step of associating further comprises the step of producing a result representing a remainder upon dividing the Previously Presented quantity by a number of network links in the plurality of network links.

27. (Previously Presented) A system for distributing data over a plurality of network links within a computer network comprising:

a processor;

a memory coupled to the processor and having locations addressable by the processor;

an operating system resident in the memory locations and executed by the processor, the operating system configured to implement a load balancing technique that defines a unit of data as a datagram, the operating system further configured to include an Internet Protocol (IP) layer that apportions the datagram into at least one fragment, the operating system further configured to include a virtual interface process that associates each datagram to a network link of the plurality of network links according to a round robin policy based at least in part on an Internet protocol (IP) identifier (ID) of each datagram; and

at least one network adapter coupled to the memory and processor to transmit the fragments of the datagram over the datagram's associated network link.

28. (Previously Presented) A system for distributing data over a plurality of network links within a computer network, comprising:

means for defining a unit of data as a datagram;

means for apportioning each datagram into at least one fragment;

means for associating each datagram to a network link of the plurality of network links according to a round robin policy based at least in part on an Internet protocol (IP) identifier (ID) of each datagram;

means for transmitting the fragments of the datagram over the datagram's associated network link.

29. (Previously Presented) A computer readable medium containing executable program instructions for execution on a processor, the executable program instructions comprising program instructions for:

defining a unit of data as a datagram;

apportioning each datagram into at least one fragment;

associating each datagram to a network link of the plurality of network links according to a round robin policy based at least in part on an Internet protocol (IP) identifier (ID) of each datagram;

transmitting the fragments of the datagram over the datagram's associated network link.

30. (Previously Presented) A method for distributing data over a plurality of network links within a computer network, comprising the steps of:

dividing a first datagram, having a first Internet protocol (IP) identifier (ID), into one or more fragments, each fragment of the first datagram associated with the first Internet protocol (IP) identifier (ID);

selecting a first network link of the plurality of network links for transmission of the one or more fragments of the first datagram;

transmitting all of the one or more fragments associated with the first IP ID over the first network link;

dividing a second datagram, having a second Internet protocol (IP) identifier (ID), into one or more fragments, each fragment of the second datagram associated with the second Internet protocol (IP) identifier (ID);

selecting a second network link of the plurality of network links for transmission of the one or more fragments of the second datagram; and

transmitting all of the one or more fragments associated with the second IP ID over the second network link.

31. (Previously Presented) The method of claim 30 wherein the first network link and the second network link are selected according to a round robin policy based on the IP ID of each datagram.

32. (Previously Presented) A system for distributing data over a plurality of network links within a computer network comprising:

a processor;

a memory coupled to the processor and having locations accessible by the processor;

an operating system resident in the memory and executed by the processor, the operating system configured to implement a load balancing technique that divides a first datagram, having a first Internet protocol (IP) identifier (ID), into one or more fragments, each fragment of the first datagram associated with the first Internet protocol (IP) identifier (ID), select a first network link of the plurality of network links for transmission of the one or more fragments of the first datagram, transmit all of the one or more fragments associated with the first IP ID over the first network link, divide a second datagram, having a second Internet protocol (IP) identifier (ID), into one or more fragments, each fragment of the second datagram associated with a second Internet protocol (IP) identifier (ID), select a second network link of the plurality of network links for transmission of the one or more fragments of the second datagram, and transmit all of the one or more fragments associated with the second IP ID over the second network link.

33. (Previously Presented) The system of claim 32 wherein the first network link and the second network link are selected according to a round robin policy based on the IP ID of each datagram.

34. (Previously Presented) The method of claim 1 wherein the step of associating further comprises apportioning data equally over the plurality of underlying links of the aggregate within the computer network.

35. (Previously Presented) The system of claim 10 wherein the virtual interface process is configured to apportion data equally over the plurality of underlying links of the aggregate within the computer network.

36. (Previously Presented) The system of claim 11 wherein the means for associating is configured to apportion data equally over the plurality of underlying links of the aggregate within the computer network.

37. (Previously Presented) The system of claim 27 wherein the round robin policy is further based, at least in part, on a number of network links in the plurality of links.

38. (Previously Presented) The system of claim 27 wherein the virtual interface process is configured to produce a result representing a remainder by dividing the IP ID by the number of network links.

39. (Previously Presented) The method of claim 38 wherein the virtual interface process is further configured to calculate the IP ID of each datagram in a sequential man-

ner and to rotate the fragments of each datagram among all the network links of the plurality of network links to thereby ensure that all fragments having the same IP ID are provided to the same network link.

40. (Previously Presented) The system of claim 27 wherein the virtual interface process is configured to logically combine the IP ID with a predetermined mask to produce a quantity, right shift the quantity by a predetermined number of places to create a Previously Presented quantity, and establish a threshold at which a group of data is forwarded to each network link of the plurality of network links.

41. (Previously Presented) The system of claim 40 wherein the virtual interface process is further configured to produce a result representing a remainder upon dividing the Previously Presented quantity by a number of network links in the plurality of network links.

42. (Previously Presented) The method of claim 30, wherein the steps of transmitting are performed in parallel.

43. (Previously Presented) A method for uniformly distributing data transmitted by a server over a plurality of underlying links within a computer network, comprising:

defining a plurality of datagram with each datagram defined as a unit of data;

apportioning each datagram into one or more fragments;

using an internet protocol (IP) identifier of each datagram to assign a link of the plurality of links to each datagram;

associating each fragment of a datagram with an IP identifier of the datagram; and

transmitting the fragments of the datagram through the assigned link with the same IP identifier.

Please add the following new claims:

44. (New) A method for uniformly distributing data transmitted by a server over a number of underlying links of an aggregate within a computer network, the method comprising the steps of:

defining a unit of data as a datagram.

assigning an Internet protocol (IP) identifier (ID) to the datagram;

apportioning the datagram into at least one fragment;

calculating a link identifier for the datagram as a function of the IP ID and the number of underlying links;

associating the fragments with the link identifier; and

transmitting the fragments over a link identified by the link identifier.

45. (New) The method of claim 44 wherein the step of calculating further comprises:

dividing the IP ID by the number of underlying links to generate a remainder; and

using the remainder as the link identifier.

46. (New) The method of claim 44 wherein the step of calculating further comprises:

performing a logical AND operation to combine the IP ID and a predetermined mask;

dividing the result of the logical AND operation by the number of underlying links to generate a remainder; and

using the remainder as the link identifier.

47. (New) The method of claim 46 wherein the predetermined mask is 0xFF80.

48. (New) The method of claim 44 wherein the step of calculating further comprises:

performing a logical AND operation to combine the IP ID and a predetermined mask;

right shifting the result of the logical AND by a predetermined number of bits;

dividing the result of right shifting by the number of underlying links to generate a remainder; and

using the remainder as the link identifier.

49. (New) The method of claim 48 wherein the predetermined mask is 0xFF80 and the predetermined number of bits is 7 bits.

50. (New) An apparatus for uniformly distributing data transmitted by a server over a number of underlying links of an aggregate within a computer network, the method comprising the steps of:

means for defining a unit of data as a datagram.

means for assigning an Internet protocol (IP) identifier (ID) to the datagram;

means for apportioning the datagram into at least one fragment;

means for calculating a link identifier for the datagram as a function of the IP ID and the number of underlying links;

means for associating the fragments with the link identifier; and

means for transmitting the fragments over a link identified by the link identifier.

51. (New) A computer readable medium containing executable program instructions for uniformly distributing data transmitted by a server over a plurality of underlying links of an aggregate within a computer network, the executable program instructions comprising program instructions for:

defining a unit of data as a datagram.

assigning an Internet protocol (IP) identifier (ID) to the datagram;

apportioning the datagram into at least one fragment;

calculating a link identifier for the datagram as a function of the IP ID and the number of underlying links;

associating the fragments with the link identifier; and

transmitting the fragments over a link identified by the link identifier.

52. (New) A method for uniformly distributing data transmitted by a server over a number of underlying links of an aggregate within a computer network, the method comprising the steps of:

selecting a datagram for transmission, the datagram having an Internet protocol (IP) identifier (ID);

selecting a link from a plurality of available links for transmitting the datagram;

breaking the datagram into fragments; and

sending all of the fragments down the selected link.

53. (New) The method of claim 52, further comprising:

calculating a link identifier for the datagram as a function of the IP ID and the number of underlying links;

associating the fragments with the link identifier; and

transmitting the fragments over a link identified by the link identifier.

54. (New) The method of claim 52, further comprising:

selecting the link by a round robin process.

55. (New) A computer network comprising:

a server; and

a plurality of clients in communication with the server over a plurality of links;

a process executing on the server, the process

(i) selecting a datagram for transmission, the datagram having an Internet protocol (IP) identifier (ID),

(ii) selecting a link from the plurality of available links for transmitting the datagram,

(iii) breaking the datagram into fragments, and

(iv) sending all of the fragments down the selected link.